A vignette in spatial areal data modelling using the **CARBayes** package.

Duncan Lee

School of Mathematics and Statistics,University of Glasgow, Glasgow, UK, G12 8QQ

March 26, 2012

1 Introduction

Data relating to a set of non-overlapping spatial units occur in many application areas, including agriculture (Besag and Higdon (1999)), ecology (Brewer and Nolan (2007)), education (Wall (2004)), epidemiology (Lee (2011)) and image analysis (Gavin and Jennison (1997)). The areal units can form a regular lattice or differ largely in both shape and size, with examples of the latter including a set of electoral wards or census tracts corresponding to a city or county. Data that occur on such areal units typically exhibit spatial correlation, with observations from units close together tending to have similar values. The spatial pattern in such data can be modelled by a hierarchical Bayesian regression model, that includes covariates and a set of random effects. The latter are included to allow for any extra variation or spatial correlation in the response data, that is not accounted for by the available covariates. The class of conditional autoregressive (CAR, Besag (1974)) models are most commonly used to capture the spatial correlation in the random effects, and are most often specified as a set of univariate conditional distributions, where the conditioning is only on the values of the random effects in geographically adjacent areal units. This geographical information is contained in a binary neighbourhood matrix W, where element w_{kj} equals one or zero depending on whether areal units (k, j) share a common border. Thus, if areas (k, j) share a common border their random effects are correlated, whereas otherwise they are conditionally independent.

In this vignette we provide a tutorial describing how to implement a number of the commonly used conditional autoregressive models in the R package **CARBayes**, as well as the recently created localised spatial smoothing method proposed by Lee and

[1]Address for correspondence - Duncan.Lee@glasgow.ac.uk.

[2]Both authors were supported by the Economic and Social Research Council [RES-000-22-4256].

Mitchell (2012). This latter approach does not force the random effects to exhibit a single level of spatial smoothness across the entire region under study, unlike the commonly used CAR models. In realistic urban contexts a single level of spatial smoothness in the response is unlikely, as there are often areas of smooth evolution, as well as locations at which step changes occur. These step changes are known as *boundaries* in the response surface, and may occur, for example, where rich and poor communities live side by side. In this context, one might expect spatial correlation to be present within each community, but not at the border where the two communities meet. The identification of these boundaries (step changes) in the response surface has a number of benefits to researchers, including the ability to detect the spatial extent of a cluster of areas that exhibit overly high or low values.

The remainder of this vignette is organised as follows. In Section 2 a brief review of the general hierarchical Bayesian model is provided, which includes the commonly used methods and a generalisation of the localised spatial smoothing approach proposed by Lee and Mitchell (2012). Then in Section 3 a tutorial of how to implement these methods in the R package **CARBayes** is presented, via a case study of alcohol disease cases in Greater Glasgow, Scotland.

2 General hierarchical Bayesian model

Recall that the region under study is split into n non-overlapping areal units, and the aim is to estimate the spatial pattern of the response across these n areas. The vector of responses is denoted by $\mathbf{Y} = (Y_1, \ldots, Y_n)$, while if a vector of offsets is required (see Section 3 for an example of where an offset could be used), it is denoted by $\mathbf{O} = (O_1, \ldots, O_n)$. In addition, there are a matrix of p covariates $X = (\mathbf{x}_1^{\mathrm{T}}, \ldots, \mathbf{x}_n^{\mathrm{T}})$, where $\mathbf{x}_k^{\mathrm{T}} = (1, x_{k2}, \ldots, x_{kp})$, the first of which corresponds to an intercept term. Finally, the vector of random effects is denoted by $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$, which are included to model any spatial correlation that remains in the residuals after the covariate effects have been accounted for. A general specification of the Bayesian hierarchical model is given by

$$\begin{aligned} Y_k|\mu_k &\sim f(y_k|\mu_k, \nu^2) \quad \text{for } k = 1, \ldots, n, \\ g(\mu_k) &= \mathbf{x}_k^{\mathrm{T}}\boldsymbol{\beta} + \phi_k + O_k, \\ \beta_i &\sim \mathrm{N}(m_i, v_i) \quad \text{for } i = 1, \ldots, p, \\ \nu_2 &\sim \mathrm{Uniform}(0, M_\nu). \end{aligned} \tag{2.1}$$

The response data Y_k come from an exponential family of distributions $f(Y_k|\mu_k, \nu^2)$, and at present the Gaussian, binomial and Poisson families are available in the **CARBayes** package. The expected value of Y_k is denoted by $\mathbb{E}[Y_k] = \mu_k$, while ν^2 is an additional scale parameter if required (for example if Y_k is Gaussian). The expected

values of the responses are related to the linear predictor via an invertible link function $g(.)$, and at present, the link functions supported are the identity (for Gaussian data), the natural log (for Poisson data) and the logit (for binomial data) functions. Independent Gaussian priors (means m_i and variances v_i) are specified for the p regression parameters β_i. A weakly informative yet proper prior distribution is assigned for the scale parameter ν^2 (if required), as little is likely to be known about its value before the analysis is conducted. We note that a common alternative would be a conjugate inverse-gamma prior for ν^2, although it is not used here because it is difficult to choose the hyperparameters so that such a prior is non-informative (for details about choosing prior distributions for variance parameters see Gelman (2006)).

The random effects model any extra variation or spatial correlation in the response data, that is not accounted for by the available covariates. However, we note that if no covariates are included in this model then the random effects represent the spatial pattern in the response surface, where as if covariates are included, they capture any structure in the residuals. The rest of this section describes a number of random effects priors for $\boldsymbol{\phi}$, which represent different types of spatial correlation structure. All of the models described can be implemented in the **CARBayes** package.

2.1 Independent random effects

If the covariate component has modelled all of the spatial correlation in the response, then only non-spatial variation (overdispersion) will remain in the residuals. In this case we denote the random effects by $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ rather than $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$ (to aid notational clarity in the following sub-sections), which leads to the altered mean model $g(\mu_k) = \mathbf{x}_k^{\mathrm{T}}\boldsymbol{\beta} + \theta_k + O_k$. The uncorrelated random effects can then be modelled as

$$\begin{aligned} \theta_k|\sigma^2 &\sim \mathrm{N}(0, \sigma^2) \qquad \text{for } k = 1, \ldots, n, \\ \sigma^2 &\sim \mathrm{Uniform}(0, M_\sigma), \end{aligned} \tag{2.2}$$

where each θ_k is independent and identically distributed with a mean of zero and a constant variance of σ^2. A diffuse uniform prior is assigned to the random effects variance σ^2, where M_σ is large enough to allow a very wide range of values.

2.2 Strong spatially correlated random effects

If the data contain spatial correlation then a common class of models for the random effects $\boldsymbol{\phi}$ are conditional autoregressive (CAR) priors, which are a type of Markov random field. As such, these models can be specified as n univariate full conditional distributions $f(\phi_k|\boldsymbol{\phi}_{-k})$, where $\boldsymbol{\phi}_{-k} = (\phi_1, \ldots, \phi_{k-1}, \phi_{k+1}, \ldots, \phi_n)$. Spatial correlation is induced into these random effects via an $n \times n$ neighbourhood matrix W, which typically contains only the numbers one and zero. In this case, if the kjth element w_{kj}

equals one then the corresponding random effects (ϕ_k, ϕ_j) are correlated. In contrast, if w_{kj} equals zero then the two random effects are conditionally independent given the values of the remaining random effects. A common specification for this neighbourhood matrix is that $w_{kj} = 1$ if areas (k, j) share a common border (and hence are spatially close), denoted $k \sim j$, otherwise w_{kj} equals zero. Other specifications are also possible, including w_{kj} equalling one if the centroids of areas (k, j) are within a threshold distance of each other, or if they are one of the q closest neighbours. Conditional on this neighbourhood information, the simplest model within the CAR class is the intrinsic autoregressive (IAR) model, which is given by

$$\begin{aligned} \phi_k|\boldsymbol{\phi}_{-k}, \tau^2, W &\sim \text{N}\left(\frac{\sum_{j=1}^n w_{kj}\phi_j}{\sum_{j=1}^n w_{kj}} \, , \, \frac{\tau^2}{\sum_{j=1}^n w_{kj}} \right) \text{for } k = 1, \ldots, n, \qquad (2.3) \\ \tau^2 &\sim \text{Uniform}(0, M_\tau). \end{aligned}$$

In this model the conditional expectation of ϕ_k is the average of the random effects in neighbouring areas (those sharing a common border and hence having w_{kj} equalling one), while the conditional variance is inversely proportional to the number of neighbouring areas. The latter recognises the fact that in the presence of strong spatial correlation, the more neighbours an area has the more information there is in the data about the value of its random effect. As before a diffuse uniform prior is specified for the variance τ^2. Finally, we note that this model can only represent strong spatial correlation, as the only hyperparameter τ^2 controls the amount of variation amongst the random effects and not the level of spatial correlation. For example, multiplying each ϕ_k by 10 will increase the variance τ^2, whilst leaving the spatial correlation structure unchanged. A further drawback of the intrinsic model is that the set of full conditional distributions above do not correspond to a proper multivariate Gaussian distribution for $\boldsymbol{\phi}$, as the corresponding precision matrix is singular.

2.3 Varying strength spatially correlated random effects

The priors in the previous two sub-sections assume that the random effects are either independent or strongly spatially correlated, which may be too restrictive for real data. Therefore a number of generalisations exist, which allow the level of spatial correlation to be estimated from the data. A number of these models are also within the class of CAR priors, and three of the most commonly used specifications can be implemented in the **CARBayes** package.

2.3.1 Besag, York and Molli (BYM) model

The BYM (Besag et al. (1991), or convolution) model is the most commonly used prior for the random effects, and combines the independent and intrinsic models described

above. The model for the mean μ_k includes two sets of random effects and is given by $g(\mu_k) = \mathbf{x}_k^{\mathrm{T}}\boldsymbol{\beta} + \theta_k + \phi_k + O_k$. The two sets of random effects (θ_k, ϕ_k) are modelled as described above, that is

$$\begin{aligned}
\phi_k|\boldsymbol{\phi}_{-k}, \tau^2, W &\sim \mathrm{N}\left(\frac{\sum_{j=1}^n w_{kj}\phi_j}{\sum_{j=1}^n w_{kj}} \, , \, \frac{\tau^2}{\sum_{j=1}^n w_{kj}} \right) \text{for } k = 1, \ldots, n, \\
\theta_k|\sigma^2 &\sim \mathrm{N}(0, \sigma^2) \qquad \text{for } k = 1, \ldots, n, \\
\tau^2 &\sim \mathrm{Uniform}(0, M_\tau), \\
\sigma^2 &\sim \mathrm{Uniform}(0, M_\sigma).
\end{aligned} \tag{2.4}$$

The advantage of this model over those described above is that it can represent varying strengths of global spatial correlation, depending on the relative sizes of the two variance parameters (τ^2, σ^2). However, the disadvantage of this model is that each data point Y_k is represented by two random effects (ϕ_k, θ_k), and as a result, only their sum is identifiable. Therefore, the estimates of the individual random effects will be unstable, and MCMC convergence may be very slow.

2.3.2 Proper model

An alternative to having two random effects is the model proposed Stern and Cressie (1999), which extends the intrinsic autoregressive model by having a correlation parameter ρ. The univariate full conditional distributions are given by

$$\begin{aligned}
\phi_k|\boldsymbol{\phi}_{-k}, \tau^2, \rho, W &\sim \mathrm{N}\left(\frac{\rho\sum_{j=1}^n w_{kj}\phi_j}{\sum_{j=1}^n w_{kj}} \, , \, \frac{\tau^2}{\sum_{j=1}^n w_{kj}} \right) \text{for } k = 1, \ldots, n, \\
\tau^2 &\sim \mathrm{Uniform}(0, M_\tau), \\
\rho &\sim \mathrm{Uniform}(0, 1).
\end{aligned} \tag{2.5}$$

This set of conditional distributions corresponds to a proper multivariate distribution for $\rho \in [0, 1)$. The conditional mean is equivalent to a weighted average of the random effects in neighbouring areas and a global mean of zero (i.e. the numerator is equal to $\rho\sum_{j=1}^n w_{kj}\phi_j + (1-\rho) \times 0$). The correlation parameter ρ allows the random effects to exhibit independence, weak or strong spatial correlation, with $\rho = 0$ corresponding to independence, while ρ close to one represents strong spatial correlation. When ρ equals one the model simplifies to the intrinsic model, while when ρ equals zero the random effects are independent and have zero mean. However, the downside of this model is that in this latter independence case the variance is still inversely proportional to the number of neighbours, even though they contain no information about ϕ_k.

2.3.3 Leroux model

An alternative to the above model which has a more natural conditional variance structure was proposed by Leroux et al. (1999), and has full conditional distributions given by

$$\begin{aligned} \phi_k|\boldsymbol{\phi}_{-k}, \tau, \rho, W &\sim \text{N}\left(\frac{\rho\sum_{j=1}^{n} w_{kj}\phi_j}{\rho\sum_{j=1}^{n} w_{kj} + 1 - \rho}\ ,\ \frac{\tau^2}{(\rho\sum_{j=1}^{n} w_{kj} + 1 - \rho)}\right), \\ \tau^2 &\sim \text{Uniform}(0, M_\tau), \\ \rho &\sim \text{Uniform}(0, 1). \end{aligned} \tag{2.6}$$

The conditional expectation is still a weighted average of the random effects in neighbouring areas and a global mean of zero (the numerator is equivalent to $\rho\sum_{j=1}^{n} w_{kj}\phi_j + (1-\rho)\times 0$). The single parameter ρ again determines the global level of spatial correlation between the random effects, with $\rho = 0$ corresponding to independence everywhere, while ρ close to one defines strong spatial correlation throughout the region. However, for this model when ρ equals zero the random effects have a constant variance, which is in keeping with the independence model.

2.4 Localised spatial smoothing models

The models described above specify a single global level of spatial smoothness or correlation for the random effects $\boldsymbol{\phi}$ across the entire study region, for example by specifying a global spatial smoothness parameter ρ. However, this is unlikely to be realistic in practice, as the spatial structure is likely to be much more localised, and include subregions of smooth evolution as well as boundaries where abrupt changes occur. The paper by Lee and Mitchell (2012) proposes a method for capturing such localised spatial structure, including the identification of boundaries in the random effects surface. Their paper was set in a disease mapping context utilising a Poisson log-normal model, and is a special case of the more general class of models that can be implemented using the **CARBayes** package. The boundaries that can be identified are in the random effects surface $\boldsymbol{\phi}$, which measures the residual pattern in disease risk after the covariate effects have been adjusted for. As a result, the model can be used in two separate ways, which effects the interpretation of any boundaries that are identified.

- If there are no covariates in the model, i.e. $\mathbf{x}_k^{\text{T}}\boldsymbol{\beta} = \beta_0$, (only an intercept term), then the mean response surface $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)$ and the random effects surface $\boldsymbol{\phi}$ have the same spatial structure, as they are related by the function $g(.)$. Therefore, any boundaries identified in the random effects surface are also boundaries in the mean response surface. It is boundaries in this latter surface that are likely to be of direct interest to researchers, in which case covariates should not be included in the model. Such a model is appropriate if the aim of the study is to identify such boundaries.

- If there are covariates in the model then the spatial structures of the random effects and response surfaces are different, so any boundaries identified may not be boundaries in the response surface. This model is appropriate if the aim of the study is to provide a realistic description of the spatial structure in the response.

The boundaries in the random effects surface are identified by modelling the set of $\{w_{kj}\}$ as random quantities if areas (k, j) share a common border, rather than assuming they are fixed at one. Conversely, if areas (k, j) do not share a common border then w_{kj} is fixed at zero. The random quantities can be modelled as binary ($w_{kj} \in \{0, 1\}$) or continuous ($w_{kj} \in [0, 1]$), and the choice will depend on the context of the problem.

- If a binary specification $w_{kj} \in \{0, 1\}$ is used, then if w_{kj} is estimated as zero (ϕ_k, ϕ_j) are conditionally independent, while if $w_{kj} = 1$ they are correlated. Thus the former corresponds to a boundary in the random effects surface between areas (k, j), while the latter corresponds to no boundary.

- In contrast, if $w_{kj} \in [0, 1]$ then the identification of boundaries is not possible, as w_{kj} can take any real number within the interval. However, unlike the previous case where $w_{kj} = 1$ or $w_{kj} = 0$, this approach does allow for varying strengths of partial correlation between (ϕ_k, ϕ_j), because its value depends on w_{kj}. Thus, this approach may be more appropriate if the goal of the analysis is to model the spatial structure in the random effects surface.

Both the binary and the continuous models for $\{w_{kj}\}$ use a small number of regression parameters $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_q)$, rather than treating each w_{kj} as a separate unknown quantity. This leads to a parsimonious covariance model for $\boldsymbol{\phi}$, whist still allowing the values of w_{kj} to be identified by the data, rather than *a-priori* by the investigator. The model is based on the hypothesis that small values of w_{kj} (hence small partial correlation between (ϕ_k, ϕ_j)) will occur between populations living in neighbouring areas that are very different, because homogeneous populations should have similar values of the reponse. Therefore, w_{kj} is modelled as a function of q non-negative dissimilarity metrics $\mathbf{z}_{kji} = (z_{ki1}, \ldots, z_{kjq})$, where

$$z_{kji} = |z_{ki} - z_{ji}|/\sigma_i, \qquad \text{for } i = 1, \ldots, q.$$

Each z_{kji} measures the absolute difference in the value of a covariate between the two areas in question, and the re-scaling by σ_i, the standard deviation of $|z_{ki} - z_{ji}|$ over all pairs of contiguous areas, improves the mixing and convergence of the MCMC algorithm. These dissimilarity metrics largely determine the spatial correlation structure across the study region, and examples could include social or physical factors. For example, social characteristics could include average income or smoking prevalence. In contrast, physical factors could include the distance between the central points of two areas, or the proportion of the boundary between the two areas that cannot be crossed on foot. Examples of the latter could include the border being a railway line, river

or motorway, and such physical barriers may prevent mixing between the two populations. Using these metrics, the proposed random effects model extends that proposed by Leroux et al. (1999) and is given by

$$\begin{aligned}
\phi_k|\boldsymbol{\phi}_{-k}, \mu, \boldsymbol{\alpha}, \tau^2 &\sim \text{N}\left(\frac{\rho\sum_{j=1}^n w_{kj}(\boldsymbol{\alpha})\phi_j}{\rho\sum_{j=1}^n w_{kj}(\boldsymbol{\alpha}) + 1 - \rho}\ ,\ \frac{\tau^2}{\rho\sum_{j=1}^n w_{kj}(\boldsymbol{\alpha}) + 1 - \rho}\right), \\
\tau^2 &\sim \text{Uniform}(0, M_\tau), \\
\rho &\sim \text{Uniform}(0, 1).
\end{aligned}$$

Here ρ determines the maximum level of spatial correlation globally, while $w_{kj}(\boldsymbol{\alpha})$ represents by how much the correlation between (ϕ_k, ϕ_j) differs from this global value. These neighbourhood relations $w_{kj}(\boldsymbol{\alpha})$ are treated as weights that range between zero and one, and if each one equals one then the global smoothing model of Leroux is obtained. The binary and continuous models for w_{kj} are given by:

Binary model

$$\begin{aligned}
w_{kj}(\boldsymbol{\alpha}) &= \begin{cases} 1 & \text{if } \exp(-\sum_{i=1}^q z_{kji}\alpha_i) \geq 0.5 \text{ and } j \sim k \\ 0 & \text{otherwise} \end{cases}, \qquad (2.7) \\
\alpha_i &\sim \text{Uniform}(0, M_i) \qquad \text{for } i = 1, \ldots, q.
\end{aligned}$$

Continuous model

$$\begin{aligned}
w_{kj}(\boldsymbol{\alpha}) &= \begin{cases} \exp(-\sum_{i=1}^q z_{kji}\alpha_i) & \text{if } j \sim k \\ 0 & \text{otherwise} \end{cases}, \qquad (2.8) \\
\alpha_i &\sim \text{Uniform}(0, M_i) \qquad \text{for } i = 1, \ldots, q.
\end{aligned}$$

For areas that are contiguous large values of (z_{kji}, α_i) result in low partial correlations, where in the binary model if $\exp(-\sum_{i=1}^q z_{kji}\alpha_i)$ is less than 0.5 then (ϕ_k, ϕ_j) are conditionally independent. The regression parameters, $\boldsymbol{\alpha}$, are constrained to be non-negative, so that the greater the dissimilarity between two areas the more likely there is to have a low (or zero) value of w_{kj}. In addition, there is no intercept term in the above model so that two areas with homogeneous populations (i.e. all $z_{kji} = 0$) have $w_{kj} = 1$. A non-informative uniform prior is assigned for each regression parameter, α_i, which corresponds to our prior ignorance about the set of w_{kj} values.

2.5 Inference

Inference for all models described in this vignette is based on Markov Chain Monte Carlo (MCMC), simulation, which is implemented in the **CARBayes** package using a combination of Metropolis and Gibbs sampling steps. The variance parameters are

Gibbs sampled from their full conditional inverse-gamma distributions, while the remaining parameters are updated using Metropolis steps with random walk proposal distributions (unless the response is Gaussian in which case Gibbs sampling can be used). Finally, to improve convergence each set of random effects is centered to have mean zero, which is implemented numerically at each iteration of the MCMC algorithm.

3 Worked example - Disease mapping

3.1 Background

Disease maps display the spatial pattern in disease risk over an extended geographical region, so that any clusters of high risk areas can be identified. For each unit only the total numbers of cases of the disease in question over a fixed time interval are available, which is compared to the number of cases expected to have occurred based on the size and demographic structure of the population living in that area. These data are typically denoted by $\mathbf{Y} = (Y_1, \ldots, Y_n)$ and $\mathbf{E} = (E_1, \ldots, E_n)$, which respectively comprise the set of observed ($\mathbf{Y}$) and expected ($\mathbf{E}$) numbers of disease cases in each of the n areal units. Using these data, the simplest measure of disease risk in area k is the ratio Y_k/E_k, known as the standardised incidence ratio (SIR) or the standardised mortality ratio (SMR). A value of one represents the null risk, where the numbers of disease cases observed and expected are the same. Values above one correspond to increased risks of disease, while values below one represent relatively healthy (low risk) areas. The general form of the model used in this context is given by

$$\begin{array}{rcll} Y_k|\mu_k & \sim & \text{Poisson}(\mu_k) & \text{for } k = 1, \ldots, n, \\ \ln(\mu_k) & = & \mathbf{x}_k^{\text{T}}\boldsymbol{\beta} + \phi_k + \ln(E_k). & \end{array} \tag{3.1}$$

This model has a mean function that can be expressed as $\mu_k = E_k R_k$, where R_k is the risk of disease in area k. Thus we have that

$$\ln(\mu_k) \;=\; \ln(E_k) + \ln(R_k) \;=\; \ln(E_k) + \mathbf{x}_k^{\text{T}}\boldsymbol{\beta} + \phi_k,$$

meaning that disease risk is represented by $R_k = \exp(\mathbf{x}_k^{\text{T}}\boldsymbol{\beta} + \phi_k)$.

3.2 Data description

We illustrate the utility of the **CARBayes** package by mapping the spatial pattern in alcohol disease in Glasgow, Scotland, between 2001 and 2004. The data are publicly available from the Scottish Neighbourhood Statistics (SNS) database (*http://www.sns.gov.uk*). The study region is the Greater Glasgow and Clyde health board, which contains the city of Glasgow in the east, and the river Clyde estuary in

Figure 1: Standardised Incidence Ratios (SIR) for alcohol admissions.

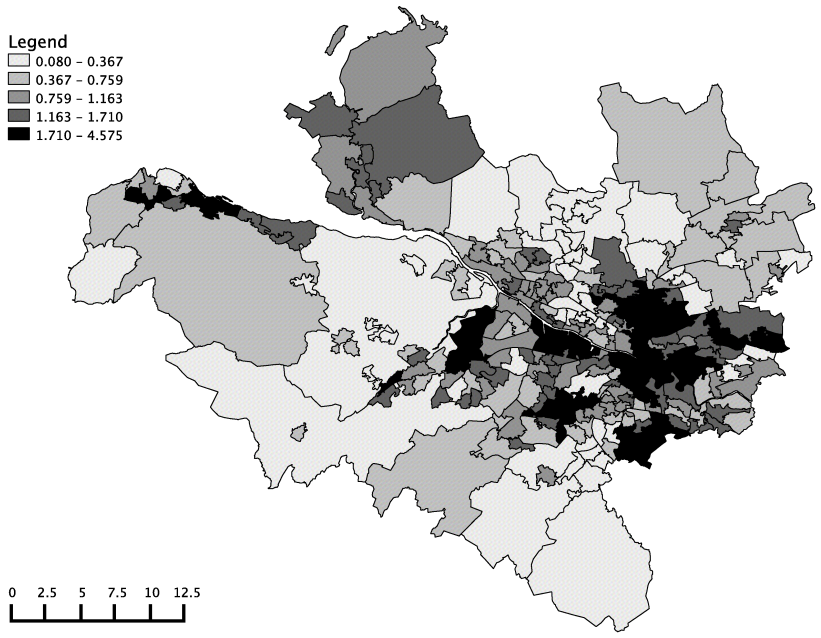

the west. Glasgow is the largest city in Scotland, with a population of around 600,000 people. It is also known to contain some of the poorest people in Europe, and has rich and poor communities that are geographically adjacent. This study region is partitioned into $n = 271$ Intermediate Geographies (IG), which were developed specifically for the distribution of small-area statistics, and have a median area of 124 hectares and a median population of 4,239.

The disease data we model are the numbers of people admitted to acute and psychiatric hospitals in each IG with a main or secondary diagnosis of alcohol related conditions, during the four year period spanning 2001 to 2004. The expected numbers of cases were calculated by external standardisation, using age and sex adjusted rates for the whole of Scotland. These rates were obtained from the Information Services Division (ISD), which is the statistical arm of the National Health Service in Scotland. The simplest measure of disease risk is the standardised incidence ratio, which is presented in Figure 1 as a choropleth map. The Figure shows that the risk of admission to hospital is highest in the heavily deprived east end of Glasgow (east of the study region), as well as along the banks of the river Clyde (the thin white line running south east).

We consider two covariates here, both of which have been shown to effect alcohol dependence and abuse rates. The first is a measure of poverty, as numerous studies (see for example Catalano et al. (1993) and Khan et al. (2002)) have shown a link between unemployment and increased alcohol dependence. The variable we consider

here is the percentage of people in each IG that are defined to be income deprived, which means they are in receipt of a combination of means-tested benefits. The second variable we consider is ethnicity, because previous studies (for example Grant et al. (2004)) have also shown a link to alcohol dependence. The only variable available to measure ethnicity is the percentage of school children in each IG from ethnic minorities (non-white). While we appreciate this variable is imperfect in many ways (e.g. it relates to children not adults, and it does not differentiate between the different ethnic groups), it is the only measure of ethnicity available for our study. The data required for this study are in two parts:

1. A data file containing the response, the expected numbers and covariates.
2. A binary $n \times n$ neighbourhood matrix containing the information on which pairs of areas are neighbours.

The following R code loads the package and reads in and formats these data

```
#### Load the package
library(CARBayes)


#### Read in the data
## Ensure the file path is correct
data <- read.csv(file="alcohol_data.csv", header=TRUE)
W <- read.table(file="W.txt", header=TRUE)
W <- as.matrix(W)

#### Format the variables
Y <- data$Y
E <- data$E
ethnic <- data$ethnic
incomdep <- data$incomedep
```

3.3 Running the independence model

The simplest random effects model assumes independence, which is summarised in Section 2.1. This model would be appropriate if, after adjusting for the covariate effects, the data contain overdispersion but not spatial correlation. This model can be run by the following code:

```
formula <- Y ~ ethnic + incomedep + offset(log(E))
burnin <- 2500
n.sample <- 5000
model1 <- poisson.independent(formula=formula, burnin=burnin, n.sample=n.sample)
```

The function can take multiple arguments, but only the *formula* one is actually required for the function to run a model. The formula argument takes the standard form specified for linear models using the *lm()* function.

The remaining arguments either have default values used by the function, or have values that are randomly generated inside the function. Default values are specified for *burnin, n.sample, blocksize.beta, blocksize.theta, prior.mean.beta, prior.var.beta, prior.max.sigma2*, which relate to the numbers of samples generated and used for inference (*burnin, n.sample*), quantities for the MCMC algorithm (*blocksize.beta, blocksize.theta*) and hyperparameters for the prior distributions (*prior.mean.beta* (m_i), *prior.var.beta* (v_i), *prior.max.sigma2* (M_σ)). Initial values for the model parameters in the MCMC algorithm are randomly generated if they are not specified, which includes *beta* ($\boldsymbol{\beta}$), *theta* ($\boldsymbol{\theta}$) and *sigma2* (σ^2).

Upon completion, the function produces the following summary table.

```
##################
#### Model fitted
##################

Likelihood model - Poisson (log link function)
Random effects model - Independent
Regression equation - Y ~ ethnic + incomedep + offset(log(E))


#############
#### Results
#############

Posterior quantiles and acceptance rates

             Median    2.5%   97.5% n.sample % accept
(Intercept) -1.1958 -1.2441 -1.1482     2500     36.1
ethnic       0.0023  0.0007  0.0038     2500     36.1
incomedep    0.0477  0.0454  0.0495     2500     36.1
sigma2       0.1414  0.1171  0.1718     2500    100.0


Acceptance rate for the random effects is 32.3%

DIC =  2334.536       p.d =  246.2667
```

The table gives posterior medians and 95% credible intervals for the regression parameters $\boldsymbol{\beta}$ and the random effects variance σ^2, as well as the number of samples on which

they are based (*n.sample* - *burnin*) and the acceptance rates (variance parameters are Gibbs sampled, hence the 100% rate). Also provided is the Deviance Information Criterion (DIC, Spiegelalter et al. (2002)) and the effective number of parameters (p.d), the former of which is a measure of how well a model fits a given data set. This can be used to compare the fit of different models to the same data set, with lower values indicating a better fitting model. The function returns a list object, which contains the following elements.

```
                Length Class   Mode
formula              3 formula call
samples.beta      7500 mcmc    numeric
samples.theta   677500 mcmc    numeric
samples.sigma2    2500 mcmc    numeric
fitted.values     1355 -none-  numeric
random.effects    1355 -none-  numeric
residuals          271 -none-  numeric
DIC                  1 -none-  numeric
p.d                  1 -none-  numeric
summary.results     20 -none-  numeric
```

- The elements *samples.beta, samples.theta, samples.sigma2* contain matrices of the MCMC samples for the parameters in the model.
- The elements *random.effects* and *fitted.values* contain summaries of the posterior distributions of the random effects $\boldsymbol{\theta}$ and the means $\boldsymbol{\mu}$ respectively for all n areas, including means, medians, standard deviations and 95% confidence intervals.
- *residuals* contains the raw residuals (i.e. $Y_k - \hat{\mu}_k$) for all n areas, where $\hat{\mu}_k$ is the posterior median.

These elements can be extracted from the model object using code such as

```
model1$fitted.values[1:5, ]
         Mean      Sd   Median     2.5%    97.5%
[1,] 153.3785 12.9791 152.4718 128.8241 180.4727
[2,]  29.9123  4.9164  29.2682  20.9858  41.0977
[3,]  49.7347  5.8273  49.6295  38.9435  60.8026
[4,]  20.8570  3.9114  21.0078  13.4794  28.4519
[5,]  81.6947  8.8957  81.5456  64.4440  98.8314
```

which extracts the summary table of the posterior distribution of the fitted values. Note, only the values from the first five areas are shown here to save space.

The estimated regression parameters are on their original scale, which is not particularly interpretable for Poisson regression models. Instead, the relative risk is typically

Table 1: Relative risks for a one standard deviation increase in each covariates value

Covariate	**Estimate**	**95% Credible interval**
Income deprivation	1.990	(1.952, 2.054)
Ethnicity	1.012	(0.981, 1.050)

specified, which measures the effect on disease risk ($\mathbf{R}$) of increasing each covariates value. For example, the relative risk for an ω increase in a covariates value measures the increase / decrease in disease risk that would occur if that covariate increased by ω. It is calculated by applying the transformation $\exp(\omega \times \cdot)$ to the posterior distribution of $\boldsymbol{\beta}$, i.e. to the median and the limits of the 95% credible interval. To illustrate, Table 1 displays the relative risk relating to a one standard deviation increase in each covariate's value. The table shows that ethnicity does not have a substantial effect on alcohol disease risk, as the 95% credible interval contains the null risk of one. In contrast, income deprivation does have a substantial effect, with a 13.8% (one standard deviation) increase in deprivation leading to a near doubling of the risk (1.990).

3.4 Running the spatial CAR models

The above model assumes the random effects are independent, which is typically an unrealistic assumption for spatial data. A number of conditional autoregressive models for capturing spatial correlation were described in the previous section, and the implementation of one of them is described here. All these models only require one additional argument compared with the independent model, namely a binary $n \times n$ neighbourhood matrix W, which contains the information on which pairs of areas are spatially close (e.g. share a common border). The remaining arguments are either identical or similar to those used for the independence model, and are described in the help files. The code to run the Leroux model is given below. As before, the remaining arguments to the function are not required for it to run.

```
#### Run the Leroux model
formula <- Y ~ ethnic + incomedep + offset(log(E))
burnin <- 2500
n.sample <- 5000
model1 <- poisson.lerouxCAR(formula=formula, W=W, burnin=burnin, n.sample=n.sample)
```

On completion the model produces a similar summary table to that of the independence model, which is given by

```
#################
#### Model fitted
#################

Likelihood model - Poisson (log link function)
Random effects model - Leroux CAR
Regression equation - Y ~ ethnic + incomedep + offset(log(E))


############
#### Results
############

Posterior quantiles and acceptance rates

             Median    2.5%   97.5% n.sample % accept
(Intercept) -1.1542 -1.2060 -1.1049     2500     33.9
ethnic      -0.0036 -0.0070 -0.0001     2500     33.9
incomedep    0.0477  0.0460  0.0498     2500     33.9
tau2         0.3355  0.2674  0.4189     2500    100.0
rho          0.8015  0.5844  0.9498     2500     80.3


Acceptance rate for the random effects is 36%

DIC =  2318.834       p.d =  236.7427
```

while the function again returns a list object similar to before.

```
summary(model1)
                Length Class   Mode
formula              3 formula call
samples.beta      7500 mcmc    numeric
samples.phi     677500 mcmc    numeric
samples.tau2      2500 mcmc    numeric
samples.rho       2500 mcmc    numeric
fitted.values     1355 -none-  numeric
random.effects    1355 -none-  numeric
residuals          271 -none-  numeric
DIC                  1 -none-  numeric
p.d                  1 -none-  numeric
summary.results     25 -none-  numeric
```

3.5 Running the localised spatial smoothing model

We illustrate the localised spatial smoothing model with binary neighbourhood relations w_{kj} and ρ fixed at 0.99 (as in Lee and Mitchell (2012)), by using it to identify boundaries in the risk surface of alcohol disease, between two areas that are geographically close but have very different risks. The identification of such boundaries is important for public health professionals, as it aids them in identifying clusters of high risk areas. As a result we do not include any covariates in the mean model (i.e. in $\mathbf{x}_k^{\mathrm{T}}\boldsymbol{\beta}$), but instead use them as dissimilarity metrics to identify possible boundaries between the random effects and risk surfaces.

The only difference between implementing this model and the other spatially correlated CAR models is the specification of the dissimilarity metrics. Each dissimilarity metric Z_i should be specified in the form of a symmetric $n \times n$ matrix, with kjth element:

$$z_{kji} = \begin{cases} |z_{ki} - z_{ji}|/\sigma_i & \text{if areas } (k,j) \text{ share a common border} \\ 0 & \text{Otherwise} \end{cases}$$

Note $z_{kki} = 0$. In this study we use our two covariates income deprivation and ethnicity as dissimilarity metrics, and as both are continuous, the corresponding matrix can be created using the following code.

```
#### Create the dissimilarity metric matrices
## Income deprivation
Z.incomedep <- as.matrix(dist(cbind(incomedep, incomedep), method="maximum",
diag=TRUE,  upper=TRUE)) * W
Z.incomedep <- Z.incomedep / sd(as.numeric(Z.incomedep[as.numeric(Z.incomedep)!=0]))

## Ethnicity
Z.ethnic <- as.matrix(dist(cbind(ethnic, ethnic), method="maximum", diag=TRUE,
upper=TRUE)) * W
Z.ethnic <- Z.ethnic / sd(as.numeric(Z.ethnic[as.numeric(Z.ethnic)!=0]))
```

For both variables, the first line creates the unscaled dissimilarity matrix, while the second scales by dividing by σ_i. Finally, these variables can be combined into a list object as follows.

```
## Combine the matrices into a list
Z <- list(Z.incomedep=Z.incomedep, Z.ethnic=Z.ethnic)
```

Once this list object is created, the model can be run using the following code.

```
## Run the model
formula <- Y ~ offset(log(E))
burnin <- 2500
n.sample <- 5000
model1 <- poisson.localisedbinaryCAR(formula=formula, rho=0.99, fix.rho=TRUE, W=W,
Z=Z, burnin=burnin, n.sample=n.sample)
```

On completion, the model produces the following summary table

```
##################
#### Model fitted
##################

Likelihood model - Poisson (log link function)
Random effects model - Localised CAR binary weights
Regression equation - Y ~ offset(log(E))
Dissimilarity metrics - Z.incomedep, Z.ethnic

#############
#### Results
#############

Posterior quantiles and acceptance rates

             Median    2.5%   97.5% n.sample % accept alpha.min
(Intercept) -0.1789 -0.1897 -0.1670     2500     35.8        NA
tau2         0.3662  0.3015  0.4506     2500    100.0        NA
Z.incomedep  0.4869  0.4781  0.5136     2500      6.0    0.1252
Z.ethnic     0.0190  0.0012  0.0694     2500      6.0    0.0992


The global spatial correlation parameter rho is fixed at 0.99

Acceptance rate for the random effects is 32.2%

DIC =  2308.99       p.d =  230.0672

The dissimilarity metrics have identified 255 borders in the study region
```

which tells us that the dissimilarity metrics have identified 255 borders in the random effects and hence the risk surface. The *alpha.min* column in the summary table gives the critical value for the parameters, below which, each dissimilarity metric has not identified any boundaries (for details see Lee and Mitchell (2012)). Thus, the table

shows that ethnicity is not responsible for identifying any boundaries, as its estimate and 95% credible interval lie below *alpha.min*. In contrast, the 95% credible interval for income deprivation lies entirely above *alpha.min*, suggesting that this variable has identified boundaries. The output from this model contains the following elements:

```
summary(model1)
                Length Class   Mode
formula              3 formula call
samples.beta      2500 mcmc    numeric
samples.phi     677500 mcmc    numeric
samples.tau2      2500 mcmc    numeric
samples.alpha     5000 mcmc    numeric
fitted.values     1355 -none-  numeric
random.effects    1355 -none-  numeric
W.posterior      73441 -none-  numeric
W.border.prob    73441 -none-  numeric
residuals          271 -none-  numeric
DIC                  1 -none-  numeric
p.d                  1 -none-  numeric
summary.results     24 -none-  numeric
```

The only additional elements compared to the other models are *W.posterior* and *W.border.prob*, which respectively give the posterior median values of w_{kj} (recall $w_{kj} = 0$ represents a boundary and $w_{kj} = 1$ represents no boundary), and the posterior probability that w_{kj} corresponds to a boundary (i.e. $w_{kj} = 0$). Finally, a map of the estimated risk surface and the locations of the boundaries is displayed in Figure 3, where the boundaries are shown in bold white lines.

References

Banerjee, S., B. Carlin, and A. Gelfand (2004). *Hierarchical Modelling and Analysis for Spatial Data* (1st ed.). Chapman and Hall.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society Series B 36*, 192–236.

Besag, J. and D. Higdon (1999). Bayesian analysis of agricultural field experiments. *Journal of the Royal Statistical Society Series B 61*, 691–746.

Besag, J., J. York, and A. Mollie (1991). Bayesian image restoration with two applications in spatial statistics. *Annals of the Institute of Statistics and Mathematics 43*, 1–59.

Brewer, M. and A. Nolan (2007). Variable smoothing in bayesian intrinsic autoregressions. *Environmetrics 18*, 841–857.

Figure 2: Estimated risk surface and the locations of boundaries.

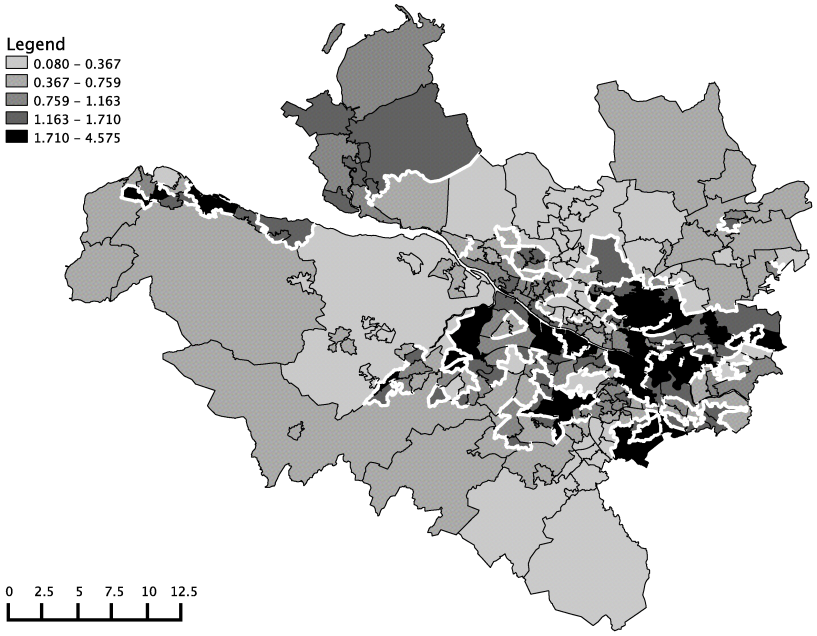

Catalano, R., D. Dooley, G. Wilson, and R. Hough (1993). Job Loss and Alcohol Abuse: A Test Using Data from the Epidemiologic Catchment Area Project. *Journal of Health and Social Behaviour 34*, 215–225.

Elliott, P., J. Wakefield, N. Best, and D. Briggs (2000). *Spatial Epidemiology: Methods and Applications* (1st ed.). Oxford University Press.

Gavin, J. and C. Jennison (1997). A subpixel Image Restoration Algorithm. *Journal of Computational and Graphical and Statistics 6*, 182–201.

Gelman, A. (2006). Prior distirbutions for variance parameters in hierarchical models. *Bayesian Analysis 1*, 515–533.

Grant, B., D. Dawson, F. Stinson, P. Chou, M. Dufour, and R. Pickering (2004). The 12-month prevalence and trends in DSM-IV alcohol abuse and dependence: United States, 19911992 and 20012002. *Drug and Alcohol Dependence 74*, 223–234.

Khan, S., R. Murray, and G. Barnes (2002). A structural equation model of the effect of poverty and unemployment on alcohol abuse. *Addictive Behaviours 27*, 405–423.

Lawson, A. (2008). *Bayesian Disease Mapping: Hierarchical Modelling in Spatial Epidemiology* (1st ed.). Chapman and Hall.

Lee, D. (2011). A comparison of conditional autoregressive model used in Bayesian disease mapping. *Spatial and Spatio-temporal Epidemiology 2*, 79–89.

Lee, D. and R. Mitchell (2012). Boundary detection in disease mapping studies. *Biostatistics to appear*, doi:10.1093/biostatistics/kxr036.

Leroux, B., X. Lei, and N. Breslow (1999). *Estimation of disease rates in small areas: A new mixed model for spatial dependence*, Chapter Statistical Models in Epidemiology, the Environment and Clinical Trials, Halloran, M and Berry, D (eds), pp. 135–178. Springer-Verlag, New York.

Spiegelalter, D., N. Best, B. Carlin, and A. Van der Linde (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society series B 64*, 583–639.

Stern, H. and N. Cressie (1999). *Inference for extremes in disease mapping*, Chapter Disease mapping and Risk Assessment for Public Health. Lawson, A and Biggeri, D and Boehning, E and Lesaffre, E and Viel, J and Bertollini, R (eds). Wiley.

Wakefield, J. (2007). Disease mapping and spatial regression with count data. *Biostatistics 8*, 158–183.

Wall, M. (2004). A close look at the spatial structure implied by the CAR and SAR models. *Journal of Statistical Planning and Inference 121*, 311–324.